www.databasejournal.com/sqletc/article.php/1469521

Back to Article

Introduction to Relational Databases
June 24, 2002

Introduction

Many web developers are self-taught, learning HTML, then moving on to a programming language such as PHP. From there, they often learn to integrate this with a database. Too few though have a good theoretical knowledge of databases. Mention foreign keys, or referential integrity, and you're met with a blank stare. Small databases can be easily designed with little database theory knowledge. But large databases can easily get out of hand when badly designed, leading to poor performance, and resulting in the whole database needing to be rebuilt later. This article is a brief introduction to the topic of relational databases, and will hopefully whet your appetite for further exploration.

The Relational Database Model

A database can be understood as a collection of related files. How those files are related depends on the model used. Early models included the hierarchical model (where files are related in a parent/child manner, with each child file having at most one parent file), and the network model (where files are related as owners and members, similar to the network model except that each member file can have more than one owner).

The relational database model was a huge step forward, as it allowed files to be related by means of a common field. In order to relate any two files, they simply need to have a common field, which makes the model extremely flexible.

Poet

Code	First Name	Surname	Age

1	Mongane	Afrika	62
2	Stephen	Serote	58
3	Tatumkhulu	Watson	29

Poem

Title	Poet
Wakening Night	1
Thrones of Darkness	2
Once	3

These two tables relate through the code field in the poet table, and the poet field in the poem table. We can see who wrote the poem 'Once' by following the relationship, and see that it was poet 3, or Tatumkhulu Watson.

In 1970, when E.F. Codd developed the model, it was thought to be hopelessly impractical, as the machines of the time could not cope with the overhead necessary to maintain the model. Of course, hardware since then has come on in huge strides, so that today even the most basic of PC's can run sophisticated relational database management systems. Together with this went the development of SQL. SQL is relatively easy to learn and allows people to quickly learn how to perform queries on a relational database. This simplicity is part of the reason that relational databases now form the majority of databases to be found.

Basic Terms

An understanding of relational databases requires an understanding of some of the basic terms.

- Data are the values stored in the database. On its own, data means very little. "43156" is an example.
- Information is data that is processed to have a meaning. For example, "43156" is the population of the town of Littlewood.
- A database is a collection of tables.
- Each table contains records, which are the horizontal rows in the table. These are also called tuples.
- Each record contains fields, which are the vertical columns of the table. These are also called attributes. An example would be a product record.
- Fields can be of many different types. There are many standard types, and each DBMS (database management system, such as Oracle or MySQL) can also have their own specific types, but generally they fall into at least three kinds - character, numeric and date. For example, a product description would be a character field, a product release date would be a date field, and a product quantity in stock would be a numeric field.
- The domain refers to the possible values each field can contain (it's sometimes called a field specification). For example, a field entitled "marital_status" may be limited to the values "Married" and "Unmarried".
- A field is said to contain a null value when it contains nothing at all. Fields can

create complexities in calculations and have consequences for data accuracy. For this reason, many fields are specifically set not to contain NULL values.

- A key is a logical way to access a record in a table. For example, in the product table, the product_id field could allow us to uniquely identify a record. A key that uniquely identifies a record is called a primary key.
- An index is a physical mechanism that improves the performance of a database. Indexes are often confused with keys. However, strictly speaking they are part of the physical structure, while keys are part of the logical structure.
- A view is a virtual table made up of a subset of the actual tables.
- A one-to-one (1:1) relationship occurs where, for each instance of table A, only one instance of table B exists, and vice-versa. For example, each vehicle registration is associated with only one engine number, and vice-versa
- A one-to-many (1:m) relationship is where, for each instance of table A, many instances of the table B exist, but for each instance of table B, only once instance of table A exists. For example, for each artist, there are many paintings. Since it is a one-to-many relationship, and not many-to-many, in this case each painting can only have been painted by one artist.
- A many to many (m:n) relationship occurs where, for each instance of table A, there are many instances of table B, and for each instance of table B, there are many instances of the table A. For example, a poetry anthology can have many authors, and each author can appear in many poetry anthologies.
- A mandatory relationship exists where, for each instance of table A, one or more instances of table B must exist. For example, for a poetry anthology to exist, there must exist at least one poem in the anthology. The reverse is not necessarily true though, as for a poem to exist, there is no need for it to appear in a poetry anthology.
- An optional relationship is where, for each instance of table A, there may exist instances of table B. For example, a poet does not necessarily have to appear in a poetry anthology. The reverse isn't necessarily true though, for example for the anthology to be listed, it must have some poets.
- Data integrity describes the accuracy, validity and consistency of data. An example of poor integrity would be where a poet's name is stored differently in two different places.
- Database normalization is a technique that helps us to reduce the occurrence of data anomalies and poor data integrity.

Page 2: Table Keys

⊙ dictionary.net

| relational database

What does relational database mean?

we found 2 entries for the meaning of **relational database**

relational database

(RDBMS - relational database management system) A <u>database</u> based on the <u>relational model</u> developed by <u>E.F. Codd</u>. A relational database allows the definition of data structures, storage and retrieval operations and <u>integrity constraints</u>. In such a database the data and relations between them are organised in <u>tables</u>. A table is a collection of <u>records</u> and each record in a table contains the same <u>fields</u>. Certain fields may be designated as <u>keys</u>, which means that searches for specific values of that field will use indexing to speed them up.

Where fields in two different tables take values from the same set, a join operation can be performed to select related records in the two tables by matching values in those fields. Often, but not always, the fields will have the same name in both tables. For example, an "orders" table might contain (customer-ID, product-code) pairs and a "products" table might contain (product-code, price) pairs so to calculate a given customer's bill you would sum the prices of all products ordered by that customer by joining on the product-code fields of the two tables. This can be extended to joining multiple tables on multiple fields. Because these relationships are only specified at retreival time, relational databases are classed as <u>dynamic database management system</u>.

The first commercial RDBMS was the Multics Relational Data Store, first sold in 1978.

<u>INGRES</u>, <u>Oracle</u>, <u>Sybase, Inc.</u>, and <u>Microsoft Access</u> are well-known database products and companies. Others include <u>PostgreSQL</u>, <u>SQL/DS</u>, <u>SYQUEST</u>, <u>IDMS</u>, and <u>RDB</u>.

["Managing Data Bases, Four Critical Factors" Michael M. Gorman, QED Information Sciences, Inc.]

["An Introduction To Database Systems" (6th ed) C. J. Date, Addison Wesley (an excellent source of detailed info)].

["An End-User's Guide to Data Base" James Martin, Prentice Hall (excellent place to begin learning about DBMS)].

(1998-10-07)

Source: The Free On-line Dictionary of Computing (13 Mar 01)

relational database

noun

a database in which relations between information items are explicitly specified as accessible attributes; "in a relational database the data are organized as a number of differently sized tables"

Source: WordNet (r) 1.7

Search for **relational database** @ <u>Ask Jeeves</u> | <u>Google</u> | <u>MSN</u> | <u>Yahoo</u>

Define relational database and 150,000 other words at dictionary.net

Definition: relational database

Search dictionary for relational database

Source: *WordNet (r) 1.7*

relational database
> n : a database in which relations between information items are
> explicitly specified as accessible attributes; "in a
> relational database the data are organized as a number of
> differently sized tables"

Source: *The Free On-line Dictionary of Computing (2003-OCT-10)*

relational database

<database> (RDBMS - relational database management system) A
database based on the relational model developed by
E.F. Codd. A relational database allows the definition of
data structures, storage and retrieval operations and
integrity constraints. In such a database the data and
relations between them are organised in tables. A table is
a collection of rows or records and each row in a table
contains the same fields. Certain fields may be designated
as keys, which means that searches for specific values of
that field will use indexing to speed them up.

Where fields in two different tables take values from the same
set, a join operation can be performed to select related
records in the two tables by matching values in those fields.
Often, but not always, the fields will have the same name in
both tables. For example, an "orders" table might contain
(customer_id, product_code) pairs and a "products" table might
contain (product_code, price) pairs so to calculate a given
customer's bill you would sum the prices of all products
ordered by that customer by joining on the product-code fields
of the two tables. This can be extended to joining multiple
tables on multiple fields. Because these relationships are
only specified at retreival time, relational databases are
classed as dynamic database management system.

The first commercial RDBMS was the Multics Relational Data
Store, first sold in 1978.

INGRES, Oracle, Sybase, Inc., Microsoft Access, and
Microsoft SQL Server are well-known database products and
companies. Others include PostgreSQL, SQL/DS, and RDB.

["Managing Data Bases, Four Critical Factors" Michael
M. Gorman, QED Information Sciences, Inc.].

["An Introduction To Database Systems" (6th ed) C. J. Date,
Addison Wesley (an excellent source of detailed info)].

["An End-User's Guide to Data Base" James Martin, Prentice
Hall (excellent place to begin learning about DBMS)].

(2002-06-10)